



Developments in Libcf and Related Software

2009 GO-ESSP Workshop,
Hamburg, Germany

Ed Hartnett

UCAR/Unidata, Oct. 6, 2009



NetCDF Java and Fortran/C/C++ Libraries

- NetCDF-Java has many features that we would like to provide to C/Fortran/C++ programmers: support for CF, advanced coordinate handling, multiple formats, etc.
- Some of this functionality is very general to scientific data, and is put in the netCDF library.
- In the C/C++/Fortran world, the CF-specific functionality is found in a separate C/Fortran library: libCF.



NetCDF-4.1 New Features

- Built-in remote data access (opendap) client.
- Can read HDF4 SD (Scientific Data) API archives.
- Can read many existing HDF5 archives.
- UDUNITS included in netCDF distribution.
- LibCF included in the netCDF distribution.
- Can do parallel I/O with classic/64-bit offset files (using parallel-netcdf library).
- Nc-config utility to help set compile flags.
- ncgen/ncdump fully handle netCDF-4 enhanced model.
- New utility: nccopy to copy netCDF files.



NetCDF Remote Data Access with OPeNDAP

- Fully C-based OPeNDAP client now built-in to netCDF.
- OPeNDAP servers can then serve many different data formats through the netCDF API.
- Developed by Dennis Heimbigner at Unidata.
- Also developing the next generation of the DAP protocol to handle the netCDF-4 enhanced model.
- A C/Fortran/C++ program can access specified subsets of remote data in a variety of formats provided by OPeNDAP servers, including Unidata's THREDDS Data Server.



More on OPeNDAP

- The DAP support is in the snapshot at essentially its final form.
- It can translate DAP to appear as either netcdf classic (the default) or netcdf-4.
- The classic translation mimics the output of libnc-dap (the older C++ library) very closely for backward compatibility.
- The opendap programmers will stop working on separate netCDF distribution.



NetCDF Reading HDF4 and HDF5 Archives

- NetCDF-4.1 C/Fortran/C++ APIs can read some existing HDF4 and HDF5 archives.
- Using this is transparent to the programmer: just open the file with a netCDF program.
- Some HDF5 types and situations are still not represented in the netCDF model (like references). These will be ignored by netCDF-4.1.
- Only HDF4 files written via the SD (Scientific Data) API are readable.



UDUNITS Distributed with NetCDF

- UDUNITS now distributed with netCDF.
- Build with `-enable-udunits`.
- NetCDF-4.1 will include UDUNITS-2.1.9.
- UDUNITS-2.1.9 has backward compatible C API, but no Fortran API yet.



Parallel I/O with parallel-netcdf Library

- Previously, netCDF-4 could do parallel I/O only with NetCDF-4/HDF5 files.
- For netCDF classic parallel I/O was only available with the parallel-netcdf library from Argonne/Northwestern U. Unfortunately it uses a different (but similar) API.
- The netCDF-4.1 release can use the parallel-netcdf library to do parallel I/O with standard netCDF API.
- The parallel-netcdf library is still used to do all I/O.

HDF Update

- HDF4 in maintain-only mode, updated once a year (latest release: HDF4.2r4).
- HDF5 under active development (latest release HDF5-1.8.3-snap2).
- HDF5 next major release: 1.10 – performance improvements, journaling, Fortran improvements.
- Much development happening for NPOESS support.



LibCF Distributed with NetCDF

- Libcf is a C/Fortran library that supports the creation of scientific data files conforming to the Climate and Forecast (CF) conventions, using the netCDF API.
- Use `--enable-libcf` to build libcf with netCDF.
- The version of libcf in the netCDF daily snapshot is **not** the latest version of libcf. For security reasons, the libcf daily development is merged into netCDF manually.



Libcf Fortran API

- The libcf Fortran API is currently under construction.
- It uses the Fortran 2003 C interoperability features, instead of the old cfortran.h method of wrapping a C library for Fortran.
- Use of fortran optional arguments means that any or all CF-related attributes can be defined in the same function call.



GRIDSPEC/libcf Motivation

- Gridspec is a PROPOSED CF standard.
- Libcf is playing a role of support for this evolving proposed standard.
- Motivation: The availability of a C/Fortran API that supports the creation and access of gridspec files will allow a wider audience of users to see whether the gridspec standard meets their model output and interoperability requirements.
- Gridspec must be explicitly turned on at libcf build time with `--enable-libcf`.



GRIDSPEC Code from GFDL

The core gridspec code was authored by Zhi Liang at GFDL.

The current libcf snapshot contains the first prototype of the GRIDSPEC library API - there will still be some changes to come in the API.

We are currently transferring Zhi Liang's development effort to the libcf copy of his code.

At Unidata the code will be included in the libcf releases, and (periodically) in netCDF/libcf releases.



Sample GRIDSPEC C Headers

```
void create_regular_lonlat_grid( int *nxbnds, int *nybnds, double *xbnds,  
                                double *ybnds,  
                                int *nlon, int *nlat, int *isc, int *iec,  
                                int *jsc, int *jec, double *x, double *y, double *dx,  
                                double *dy, double *area, double *angle_dx, char *center_cell );  
  
void create_simple_cartesian_grid( double *xbnds, double *ybnds, int *nlon, int *nlat,  
                                   double *simple_dx, double *simple_dy, int *isc, int *iec,  
                                   int *jsc, int *jec, double *x, double *y,  
                                   double *dx, double *dy, double *area, double  
                                   *angle_dx, char *center_cell );  
  
void create_grid_from_file( char *file, int *nlon, int *nlat, double *x, double *y,  
                            double *dx, double *dy,  
                            double *area, double *angle_dx );
```

etc...



GRIDSPEC Automated Testing

In line with the usual Unidata development style, GRIDSPEC code will be subjected to regular, automated testing.

Daily snapshot builds will be available from Unidata FTP site.

Output of libcf testing is distributed via email to developers.

A reference CDL file (generated by the current command line tool, and inspected) is used to test various GRISPEC features.

Some testing is already taking place, but much more remains to be added.



GRIDSPEC Documentation

Documentation for the C API is available in the libcf documentation.

Current documentation is based on the command line tool documentation.



The Future of GRISPEC/libcf

Full integration of Zhi Liang's development effort is underway. When that is complete, his latest code changes will be merged, and the libcf/gridspec codebase will become the home of the gridspec code.

Other developers are invited to collaborate in this development process. Email ed@unidata.ucar.edu if you are interested in contributing to GRIDSPEC development.

Role of Beta-Testers and Early Adopters in NetCDF

- There are three programmers on the netCDF C/Fortran libraries. We also do all documentation and support.
- Hundreds of thousands of automated tests and daily snapshot releases give us the confidence to develop rapidly.
- Development works best when interested users try out important features in the snapshot release, and communicate with the developers.

